arXiv:1703.02251v1 [math.AG] 7 Mar 2017

THE MAXIMUM LIKELIHOOD DEGREE OF TORIC VARIETIES

CARLOS AMÉNDOLA, NATHAN BLISS, ISAAC BURKE, COURTNEY R. GIBBONS, MARTIN HELMER, SERKAN HOŞTEN, EVAN D. NASH, JOSE ISRAEL RODRIGUEZ, DANIEL SMOLKIN

Abstract. We study the maximum likelihood degree (ML degree) of toric varieties, known as discrete exponential models in statistics. By introducing scaling coefficients to the monomial parameterization of the toric variety, one can change the ML degree. We show that the ML degree is equal to the degree of the toric variety for generic scalings, while it drops if and only if the scaling vector is in the locus of the principal A-determinant. We also illustrate how to compute the ML estimate of a toric variety numerically via homotopy continuation from a scaled toric variety with low ML degree. Throughout, we include examples motivated by algebraic geometry and statistics. We compute the ML degree of rational normal scrolls and a large class of Veronese-type varieties. In addition, we investigate the ML degree of scaled Segre varieties, hierarchical loglinear models, and graphical models.

1. Introduction

Maximum likelihood estimation is a fundamental optimization problem in statistical inference. The maximum likelihood degree of an algebraic statistical model was introduced in [CHKS06] and [HKS05] to study the geometry and complexity of this optimization problem. Here we study the maximum likelihood degree of toric varieties.

Let $V \subset \mathbb{P}^{n-1}$ be a projective variety over $\mathbb{C}$. The homogeneous coordinates of $\mathbb{P}^{n-1}$ will be denoted by $[p_1 : \cdots : p_n]$. The affine open subset of V where $p_1 + \ldots + p_n \neq 0$ is identified with the set of points in V where $p_1 + \ldots + p_n = 1$. We can view the points $p \in V$ satisfying

$$p_1, \ldots, p_n \geq 0 \text{ and } p_1 + \ldots + p_n = 1$$

as a family of probability distributions of a discrete random variable X where

$$p_i = \text{prob}(X = i) \quad \text{for } i = 1, \ldots, n.$$

After observing X in N instances $X_1, \ldots, X_N$, we record the data vector $u = (u_1, \ldots, u_n) \in \mathbb{Z}^n$ where u_i is the number instances where $X_j = i$, i.e.,

$$u_i = |\{j : X_j = i\}| \text{ and } N = u_1 + \cdots + u_n.$$

The likelihood function is the rational function

$$\ell_u(p) = \frac{p_1^{u_1} \cdots p_n^{u_n}}{(p_1 + \cdots + p_n)^{u_1 + \cdots + u_n}},$$

and one seeks to find a probability distribution $\hat{p} = (\hat{p}_1, \ldots, \hat{p}_n)$ in V which maximizes ℓ_u. Such a probability distribution $\hat{p}$ is a maximum likelihood estimate, and $\hat{p}$ can be identified by computing all critical points of ℓ_u on V.

This project was started during the AMS Mathematics Research Communities 2016 program in Algebraic Statistics held in Snowbird, Utah. The authors were part of the Likelihood Geometry group led by Serkan Hoşten and Jose Rodriguez. This material is based upon work supported by the National Science Foundation under Grant Number DMS 1321794.

We let V_{sing} be the singular locus of V and $V_{\text{reg}} = V \setminus V_{\text{sing}}$ be the regular locus. We also define $\mathcal{H}$ as the hypersurface that is the union of the coordinate hyperplanes and the hyperplane defined by $p_1 + \cdots + p_n = 0$.

Definition 1.1. The *maximum likelihood degree* of V, denoted $\text{mldeg}(V)$, is the number of complex critical points of the likelihood function ℓ_u on $V_{\text{reg}} \setminus \mathcal{H}$ for generic vectors u.

In [CHKS06] and [HKS05], it was shown that $\text{mldeg}(V)$ is well-defined. We refer the reader to [DSS09, Huh13, HRS14] for a glimpse of the growing body of work on aspects of $\text{mldeg}(V)$. For an excellent recent survey we recommend [HS14].

In most applications, V is the Zariski closure of the image of a polynomial map, also known as a parametric statistical model. A widely used subclass of such models is composed of those given by a monomial parametrization. In algebraic geometry, they are known as toric varieties [Ful93, CLS11], and in probability theory and statistics they are known as discrete exponential families [BD76]. In particular, hierarchical loglinear models [BFH07] and undirected graphical models on discrete random variables [Lau96] present examples in contemporary applications. Theorem 3.1 is our main result and is restated below.

Theorem (Main result)**.** *Let $V^c \subset \mathbb{P}^{n-1}$ be the projective variety defined by the monomial parametrization $\psi^c : (\mathbb{C}^*)^d \longrightarrow (\mathbb{C}^*)^n$ where*

$$\psi^c(s, \theta_1, \theta_2, \ldots, \theta_{d-1}) = (c_1 s\theta^{a_1}, c_2 s\theta^{a_2}, \ldots, c_n s\theta^{a_n}),$$

and $c \in (\mathbb{C}^*)^n$ is fixed. Then $\text{mldeg}(V^c) < \deg(V)$ if and only if c is in the principal A-determinant of the toric variety $V = V^{(1,\ldots,1)}$.

To set the stage, we start with an example chosen from the theory of graphical models. This is the smallest example of a graphical model that is not decomposable. Models that are decomposable have a unique maximum likelihood estimate that is a rational function of the data vector u. This is equivalent to $\text{mldeg}(V) = 1$ [GMS06].

Example 1.2 (Binary 4-Cycle)**.** This example is from a study to determine risk factors for coronary heart disease based on data collected in Czechoslovakia [R+81, EH85]. Six different factors affecting coronary heart disease were recorded for a sample of 1841 workers employed in the Czech automotive industry. We will only use the following four factors: whether they smoked (S), whether their systolic blood pressure was less than 140 mm (B), whether there was a family history of coronary heart disease (H), and whether the ratio of beta to alpha lipoproteins was less than 3 (L). Each is a binary variable, and the joint random variable $X = (S, B, H, L)$ has a state space of cardinality 16. The data set is summarized in Table 1.

We will use the graphical model known as the 4-cycle (Figure 1). If we set

$$p_{ijk\ell} = \text{prob}(S = i, B = j, H = k, L = \ell)$$

the family of probability distributions for X which factor according to this model can be described by the following monomial parametrization: let $a_{ij}, b_{jk}, c_{k\ell}, d_{i\ell}$ be parameters for $i, j, k, \ell \in \{0, 1\}$ and let $p_{ijk\ell} = a_{ij}b_{jk}c_{k\ell}d_{i\ell}$. The Zariski closure V of the image of this parametrization is a toric variety. Every probability distribution for X that lies on this toric variety satisfies certain independence statements which can be read from the graph. For instance, S and H are independent given B and L. Similarly, B and L are independent given

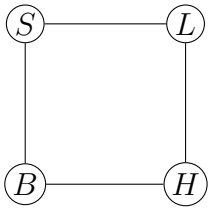

FIGURE 1. The 4-cycle.

H	L	B	S: no	S: yes
neg	< 3	< 140	297	275
		≥ 140	231	121
	≥ 3	< 140	150	191
		≥ 140	155	161
pos	< 3	< 140	36	37
		≥ 140	34	30
	≥ 3	< 140	32	36
		≥ 140	26	29

TABLE 1. Worker data.

S and H. The maximum likelihood estimate given by the worker data is

$$\begin{aligned}\hat{p} = (\, & 0.15293342,\ 0.089760679,\ 0.021266977,\ 0.015778191,\\ & 0.12976986,\ 0.076165372,\ 0.020853199,\ 0.015471205,\\ & 0.13533793,\ 0.11789409,\ 0.018820142,\ 0.0207235,\\ & 0.083859917,\ 0.073051125,\ 0.01347576,\ 0.014838619\,).\end{aligned}$$

The degree of V is 64 and $\mathrm{mldeg}(V) = 13$. This was first computed in [GMS06, p. 1484] where the question of explaining the fact that $\mathrm{mldeg}(V) = 13$ was raised.

The article is organized as follows. In Section 2 we define scaled toric varieties, introduce maximum likelihood degrees, and recall discriminants. In Section 3 we prove our main theorem (Theorem 3.1). In Sections 4-7 we compute the ML degree of rational normal scrolls and a large class of Veronese-type varieties. We investigate the ML degree of scaled Segre varieties, hierarchical loglinear models, and graphical models. In Section 8, we show that maximum likelihood estimates can be tracked between different scalings of a toric variety via homotopy continuation. We illustrate with computational experiments that homotopy continuation can be a better numerical method than iterative proportional scaling [DR72] for computing the maximum likelihood estimate of a discrete exponential model.

2. Preliminaries

In this section we introduce scaled toric varieties and their likelihood equations. We show that the ML degree of a scaled toric variety is bounded by the degree of the variety. Next, we present two classical objects in toric geometry: the A-discriminant and the principal A-determinant [GKZ94]. The principal A-determinant plays a crucial role in our main result Theorem 3.1.

2.1. **Scaled toric varieties.** In this paper, we study the maximum likelihood degree of projective toric varieties. Let A be a $(d-1) \times n$ matrix with columns $a_1, \ldots, a_n \in \mathbb{Z}^{d-1}$. The convex hull of the lattice points a_i defines a polytope $Q = \mathrm{conv}(A)$.

Definition 2.1. A toric variety scaled by $c \in (\mathbb{C}^*)^n$, denoted V^c, is defined by the following map $\psi^c : (\mathbb{C}^*)^d \longrightarrow (\mathbb{C}^*)^n$ where A is a full rank $(d-1) \times n$ integer matrix:

$$\psi^c(s, \theta_1, \theta_2, \ldots, \theta_{d-1}) = (c_1 s\theta^{a_1}, c_2 s\theta^{a_2}, \ldots, c_n s\theta^{a_n}). \tag{1}$$

Here V^c is the projective variety in $\mathbb{P}^{n-1}$ corresponding to the affine cone that is the closure of $\psi^c((\mathbb{C}^*)^d)$ in $\mathbb{C}^n$.

Different choices of c give different embeddings of isomorphic varieties V^c in $\mathbb{P}^{n-1}$; their ML degrees may also differ. When $c = (1, 1, \ldots, 1, 1)$, V^c is the standard toric variety of A as in [CLS11].

Example 2.2. Consider $\psi^c : (\mathbb{C}^*)^3 \longrightarrow (\mathbb{C}^*)^6$ given by

$$(s, t, u) \mapsto (s, 2st, st^2, 2su, 2stu, su^2).$$

This gives the Veronese embedding of $\mathbb{P}^2$ into $\mathbb{P}^5$ scaled by $c = (1, 2, 1, 2, 2, 1)$. The usual Veronese embedding has ML degree four whereas this scaled version has ML degree one.

2.2. **Likelihood equations.** We begin our discussion with the critical equations whose solutions are the critical points of the likelihood function for a toric variety. These equations are called likelihood equations.

Let A be a $(d-1) \times n$ matrix as in Definition 2.1. We let

$$f = \sum_{i=1}^{n} c_i \theta^{a_i} = c_1 \theta^{a_1} + \cdots + c_n \theta^{a_n}.$$

Definition 2.3. Let $u = (u_1, \ldots, u_n)$ and $u_+ = \sum_i u_i$. Using the method of Lagrange multipliers, we obtain the likelihood equations for the variety V^c:

$$1 = sf$$
$$(Au)_i = u_+ s\theta_i \frac{\partial f}{\partial \theta_i} \text{ for } i = 1, \ldots, d-1.$$

In other words,

$$\begin{aligned} 1 &= sf \\ (Au)_1 &= u_+ s\theta_1 \tfrac{\partial f}{\partial \theta_1} \\ (Au)_2 &= u_+ s\theta_2 \tfrac{\partial f}{\partial \theta_2} \\ &\vdots \\ (Au)_{d-1} &= u_+ s\theta_{d-1} \tfrac{\partial f}{\partial \theta_{d-1}} \end{aligned} \tag{2}$$

We remark that the solutions to the above system are exactly the solutions to

$$\begin{aligned} (Au)_1 f &= u_+ \theta_1 \tfrac{\partial f}{\partial \theta_1} \\ (Au)_2 f &= u_+ \theta_2 \tfrac{\partial f}{\partial \theta_2} \\ &\vdots \\ (Au)_{d-1} f &= u_+ \theta_{d-1} \tfrac{\partial f}{\partial \theta_{d-1}} \end{aligned} \tag{3}$$

where $f \neq 0$.

We formulate the ML degree of the scaled toric variety V^c using the coordinates $p_1, \ldots, p_n$ as well. Note that $f = c_1\theta^{a_1} + \cdots + c_n\theta^{a_n}$ induces a linear form $L_f(p) = c_1p_1 + \cdots + c_np_n$ on the toric variety V. Similarly, $\theta_i \frac{\partial f}{\partial \theta_i}$ induce corresponding linear forms $L_i(p)$ on V for $i = 1, \ldots, d-1$, as the monomials appearing in $\theta_i \frac{\partial f}{\partial \theta_i}$ are a subset of the monomials appearing in f. With this we immediately get the following.

Proposition 2.4. *The ML degree of V^c is the number of solutions $p \in V \setminus V(p_1 \cdots p_n L_f(p))$ to*

$$\begin{array}{rcl} (Au)_1 L_f(p) & = & u_+ L_1(p) \\ (Au)_2 L_f(p) & = & u_+ L_2(p) \\ & \vdots & \\ (Au)_{d-1} L_f(p) & = & u_+ L_{d-1}(p) \end{array} \tag{4}$$

for generic vectors u.

Proof. Definition 1.1 combined with (3) give the result by noting the fact that the singularities of any toric variety are contained in the union of coordinate hyperplanes. □

The following result was first proved in [HS14, Theorem 3.2].

Corollary 2.5. *For $c \in (\mathbb{C}^*)^n$, the ML degree of V^c is at most $\deg(V)$. Moreover, for generic c, $\mathrm{mldeg}(V^c) = \deg(V)$.*

Proof. For generic u, the linear equations (4) define a linear subspace of codimension $d-1$. By Bertini's theorem, the intersection of this linear subspace with V is generically transverse. Now by Bezout's theorem [Ful98, Propositon 8.4], the sum of the degrees of the components of the intersection is equal to the degree of V. Therefore, the degree of the zero dimensional piece of this intersection is at most $\deg(V)$. The ML degree of V^c is obtained by removing those solutions in $V(p_1 \cdots p_n L_f(p))$, hence $\mathrm{mldeg}(V^c) \leq \deg(V)$. For generic c, the linear subspace is a generic linear subspace of codimension $d-1$. Not only will the intersection contain exactly $\deg(V)$ points, these points will not be in $V(p_1 \cdots p_n L_f(p))$. Hence, in this case, $\mathrm{mldeg}(V^c) = \deg(V)$. □

We will denote the set of solutions in V to (4) by $\mathcal{L}'_{c,u}$. Similarly, we will denote the set of solutions to (4) in $V(p_1 \cdots p_n L_f(p))$ by $\mathcal{L}_{c,u}$. To close this subsection, we would like to point out a classical result both in probability theory and toric geometry. It is sometimes known as Birch's theorem (see e.g. [Lau96]).

Theorem 2.6 (Birch's Theorem)**.** *There is a unique positive point in $\mathcal{L}_{c,u}(p)$, namely, there is a unique positive solution to (4) in $V \setminus V(p_1 \cdots p_n L_f(p))$ if the scaling vector c and the data vector u are positive vectors.*

2.3. **A-discriminant and principal A-determinant.** In this subsection we recall the classical definitions of discriminants and resultants.

Definition 2.7. To any matrix A as above, we can associate the variety ∇_A, defined by

$$\nabla_A = \overline{\left\{ c \in (\mathbb{C}^*)^n \mid \exists \theta \in (\mathbb{C}^*)^{d-1} \text{ such that } f(\theta) = \frac{\partial f}{\partial \theta_i}(\theta) = 0 \text{ for all } i \right\}}.$$

If ∇_A has codimension one in $(\mathbb{C}^*)^n$, then the *A-discriminant*, denoted $\Delta_A(f)$, is defined to be the irreducible polynomial that vanishes on ∇_A. This polynomial is unique up to multiplication by a scalar.

As in the previous subsection, to $f = \sum_{i=1}^n c_i \theta^{a_i}$ where $c = (c_1, \ldots, c_n) \in (\mathbb{C}^*)^n$ we associate the linear form $L_f(p) = \sum_{i=1}^n c_i p_i$.

Definition 2.8. Given $f_1, \ldots, f_d$ where $f_j = \sum_{i=1}^n c_{i,j}\theta^{a_i}$ with coefficients $c_{i,j} \in \mathbb{C}^*$ there is a unique irreducible polynomial $R_A(f_1, \ldots, f_d)$ in the coefficients $c_{i,j}$, so that $R_A(f_1, \ldots, f_d) = 0$ at $(c_{i,j})$ if and only if $L_{f_1} = L_{f_2} = \cdots = L_{f_d} = 0$ on the toric variety V [GKZ94, Section 2, Chapter 8]. This polynomial is called the *A-resultant* of $f_1, \ldots, f_d$.

Definition 2.9. For $f = \sum_{i=1}^n c_i\theta^{a_i}$, we define the *principal A-determinant* as

$$E_A(f) \;=\; R_A\left(f, \theta_1\frac{\partial f}{\partial \theta_1}, \ldots, \theta_{d-1}\frac{\partial f}{\partial \theta_{d-1}}\right).$$

This is a polynomial in the coefficients c_i of f.

In [GKZ94], the principal A-determinant $E_A(f)$ is related to the polytope Q of the toric variety V. More precisely, when the toric variety V is smooth, and A is the matrix whose columns correspond to the lattice points in Q, the principal A-determinant is

$$E_A(f) \;=\; \prod_{\Gamma \text{ face of } Q} \Delta_{\Gamma\cap A} \tag{5}$$

where the product is taken over all nonempty faces $\Gamma \subset Q$ including Q itself [GKZ94, Theorem 1.2, Chapter 10]. Here, $\Gamma \cap A$ is the matrix whose columns correspond to the lattice points contained in Γ. When V is not smooth, the radical of its principal A-determinant is the polynomial given in (5). The zero locus of $E_A(f)$ in $\mathbb{C}^n$ will be denoted by Σ_A.

3. ML degree and principal A-determinant

We first state and prove our main theorem.

Theorem 3.1. *Let $A = (a_1 \ldots a_n)$ be a full rank $(d-1) \times n$ integer matrix, and let $V^c \subset \mathbb{P}^{n-1}$ be the scaled toric variety defined by the monomial parametrization given by (1) where $c \in (\mathbb{C}^*)^n$ is fixed. Then* $\mathrm{mldeg}(V^c) < \deg(V^{(1,1,\ldots,1,1)})$ *if and only if $c \in \Sigma_A$.*

Proof. Let $V = V^{(1,1,\ldots,1,1)}$. Let $c \in (\mathbb{C}^*)^n \cap \Sigma_A$. Then there exists $p \in V$ such that $L_f(p) = L_1(p) = \cdots = L_{d-1}(p) = 0$. Such a p is a solution in $\mathcal{L}'_{c,u}$ but not a solution in $\mathcal{L}_{c,u}$. Since the degree of $\mathcal{L}'_{c,u}$ is equal to $\deg(V)$, and since $\mathrm{mldeg}(V^c)$ is equal to the degree of $\mathcal{L}_{c,u}$ we see that $\mathrm{mldeg}(V^c) < \deg(V)$.

Conversely, suppose $\mathrm{mldeg}(V^c) < \deg(V)$. There are two ways this can happen. On the one hand, there is $p \in V$ so that $L_f(p) = 0$ which implies that $L_1(p) = \cdots = L_{d-1}(p) = 0$. This means $c \in \Sigma_A$. On the other hand, there is a solution $p \in V$ in $\mathcal{L}'_{c,u}$ where some $p_i = 0$. The coordinates of p that are zero cannot be arbitrary. The support of p, that is $\{i : p_i \neq 0\}$, is in bijection with the columns a_i of A where the convex hull of these columns is a face of $Q = \mathrm{conv}(A)$. Let Γ be the corresponding face. Without loss of generality we will assume that $\Gamma \cap A = \{a_1, \ldots, a_k\}$. Let $e < d-1$ be the dimension of Γ. Since $p_{k+1} = \cdots = p_n = 0$, the point $p_\Gamma = [p_1 : \cdots : p_k]$ is in the toric variety V_Γ defined by $\Gamma \cap A$. Moreover, $L_f(p) = L_{f_\Gamma}(p_1, \ldots, p_k)$ where $L_{f_\Gamma} = \sum_{i=1}^k c_i p_i$ is the linear form associated to the polynomial f_Γ whose support is Γ. Similarly, $L_i(p) = L'_i(p_1, \ldots, p_k)$ where L'_i is the linear

form associated to $\theta_i \frac{\partial f_\Gamma}{\partial \theta_i}$. If $L_f(p) = L_{f_\Gamma}(p_1, \ldots, p_k) \neq 0$, then p_Γ is a solution to

$$\begin{array}{rcl} 1 & = & L_{f_\Gamma} \\ (Au)_1 & = & u_+ L'_1 \\ (Au)_2 & = & u_+ L'_2 \\ & \vdots & \\ (Au)_{d-1} & = & u_+ L'_{d-1}. \end{array}$$

But this is a $d \times k$ linear system of rank $e < d$. Because u is generic, this system does not have a solution. Therefore the only way $p = [p_1 : \cdots : p_k : 0 : \cdots : 0]$ can be a solution in $\mathcal{L}'_{c,u}$ is with $L_{f_\Gamma}(p_1, \ldots, p_k) = L_f(p) = 0$. We again conclude that $c \in \Sigma_A$. □

Remark 3.2. A natural question one can ask is under what conditions a general projective variety has ML degree one. This question was answered in [Huh14].

3.1. **Toric Hypersurfaces.** Let $A = (a_1 \ \ldots \ a_{d+1})$ be a $(d-1) \times (d+1)$ integer matrix of rank $d-1$. In this case, the toric variety V and each scaled toric variety V^c is a hypersurface generated by a single polynomial. This polynomial can be computed as follows. We let A' be the matrix obtained by appending a row of 1's to the matrix A. The kernel of A' is generated by an integer vector $w = (w_1, \ldots, w_{d+1})$ where $\gcd(w_1, \ldots, w_{d+1}) = 1$. Without loss of generality we assume that $w_1, \ldots, w_\ell > 0$ and $w_{\ell+1}, \ldots, w_{d+1} \leq 0$. Then the hypersurface is defined by

$$p_1^{w_1} p_2^{w_2} \cdots p_\ell^{w_\ell} - p_{\ell+1}^{-w_{\ell+1}} \cdots p_{d+1}^{-w_{d+1}}.$$

We say A' and A are in general position if no d columns of A' lie on a hyperplane. The matrix A' is in general position if and only if w has full support, i.e. no w_i is 0. This is also equivalent to $Q = \operatorname{conv}(A)$ being a simplicial polytope with exactly d of the columns of A on each facet of Q. Recall that $f = \sum_{i=1}^{d+1} c_i \theta_i^{a_i}$. In this setting, the A-discriminant Δ_A and the principal A-determinant $E_A(f)$ can be calculated directly.

Proposition 3.3. [GKZ94, Chapter 9, Proposition 1.8] *If A is in general position then*

$$\Delta_A = \left(w_{\ell+1}^{-w_{\ell+1}} \cdots w_{d+1}^{-w_{d+1}}\right) c_1^{w_1} \cdots c_\ell^{w_\ell} - \left(w_1^{w_1} \cdots w_\ell^{w_\ell}\right) c_{\ell+1}^{-w_{\ell+1}} \cdots c_{d+1}^{-w_{d+1}}.$$

This is also the principal A-determinant of this toric hypersurface up to a monomial factor.

Corollary 3.4. *If A is in general position,* $\operatorname{mldeg}(V^c) < \deg(V)$ *if and only if*

$$c \in V\left(\left(w_{\ell+1}^{-w_{\ell+1}} \cdots w_{d+1}^{-w_{d+1}}\right) c_1^{w_1} \cdots c_\ell^{w_\ell} - \left(w_1^{w_1} \cdots w_\ell^{w_\ell}\right) c_{\ell+1}^{-w_{\ell+1}} \cdots c_{d+1}^{-w_{d+1}}\right)$$

where $p_1^{w_1} p_2^{w_2} \cdots p_\ell^{w_\ell} - p_{\ell+1}^{-w_{\ell+1}} \cdots p_{d+1}^{-w_{d+1}}$ is the polynomial defining the hypersurface V.

4. Rational Normal Scrolls

In this section, we characterize the ML degree of a rational normal scroll. A *rational normal scroll* is a toric variety associated to a $(d-1) \times n$ matrix A of the following form:

$$A = \begin{pmatrix} 1 & \cdots & 1 & 0 & \cdots & 0 & \cdots & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & \cdots & 0 & 1 & \cdots & 1 & \cdots & 0 & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & & & & & & & & & & & & \vdots \\ 0 & \cdots & 0 & 0 & \cdots & 0 & \cdots & 1 & \cdots & 1 & 0 & \cdots & 0 \\ 0 & 1 \cdots & n_1 & 0 & 1 \cdots & n_2 & \cdots & 0 & 1 \cdots & n_{d-2} & 0 & 1 \cdots & n_{d-1} \end{pmatrix}. \tag{6}$$

Here, $n = d - 1 + n_1 + n_2 + \cdots + n_{d-1}$. This description is due to Petrović [Pet08].

We index the data vector and scaling vector as (u_{ij}) and (c_{ij}) respectively. Define g_i to be the following univariate polynomials:

$$g_i = c_{i0} + c_{i1}\theta_i + c_{i2}\theta_i^2 + \cdots + c_{in_i}\theta_i^{n_i}, \text{ for } i = 1, 2, \ldots, d-1.$$

Note that for $i \neq d-1$, g_i is the ith partial derivative of f where

$$f = (\theta_1 g_1 + \theta_2 g_2 + \cdots + \theta_{d-2} g_{d-2}) + g_{d-1}.$$

We now state and prove our result on the ML degree of the rational normal scroll.

Theorem 4.1. *The ML degree of a rational normal scroll given by the matrix A as in (6) is equal to the number of distinct roots of $g_1 g_2 \cdots g_{d-1}$.*

Proof. For each i, we use g_i' to denote the derivative of g_i with respect to θ_{d-1}, and we write $u_{(d-1)+} = \sum_{j=0}^{n_{d-1}} u_{(d-1)j}$. With this set up, the last equation in (3) is

$$(Au)_{d-1} f = u_+ \theta_{d-1} \left(g_{d-1}' + \sum_{i=1}^{d-2} \theta_i g_i' \right).$$

Multiply both sides of this equation by $g_1 g_2 \cdots g_{d-1}$. Then substituting $(Au)_i f = u_+ \theta_i g_i$ for each $1 \leq i \leq d-2$ in the appropriate summand on the right-hand side gives:

$$(Au)_{d-1} g_1 \cdots g_{d-1} = \theta_{d-1} \left(u_{(d-1)+} g_1 \cdots g_{d-2} g_{d-1}' + \sum_{i=1}^{d-2} (Au)_i g_1 \cdots g_{i-1} g_i' g_{i+1} \cdots g_{d-1} \right). \tag{7}$$

For generic (u_{ij}), the solutions to (3) correspond to the solutions of (7) where $g_i \neq 0$ for every i.

Suppose that $\theta_{d-1} = \alpha$ is a repeated root of $g_1 g_2 \cdots g_{d-1}$. Then either α is a shared root of g_i and g_j for $i \neq j$ or α is a repeated root of a single g_i. In the first case, $\theta_{d-1} - \alpha$ divides the right-hand side of (7) since each summand has either g_i or g_j as a factor. In the second case, $\theta_{d-1} - \alpha$ divides both g_i and g_i' and hence the right-hand side of (7). Factoring out all such repeated roots results in an equation whose degree equals the number of distinct roots of $g_1 \cdots g_{d-1}$. This shows that the ML degree is at most the number of the distinct roots; we must still argue that no more roots can be factored out.

Suppose $g_1 g_2 \cdots g_{d-1}$ has no repeated roots and let $\theta_{d-1} = \beta$ be a root of an arbitrary g_i. Then $\theta_{d-1} - \beta$ does not divide $g_1 \cdots g_{i-1} g_i' g_{i+1} \cdots g_{d-1}$ but does divide every other summand on the right-hand side, so it does not divide the sum as a whole. This follows because for generic data vectors (u_{ij}), no coefficient on the right-hand side vanishes. Thus, if the repeated roots of $g_1 g_2 \cdots g_{d-1}$ have been removed, the degree of the polynomial and hence the ML degree will not drop any further. □

We finish this section with examples.

Example 4.2. The simplest example of a rational normal scroll is a rational normal curve with $A = (0\,1\,\ldots\,n)$. We denote the scaled rational normal curve of degree n by $\mathcal{C}_n^c$. The ML degree of the scaled rational normal curve $\mathcal{C}_n^c$ is equal to the number of distinct roots of the polynomial

$$c_0 + c_1\theta + c_2\theta^2 + \cdots + c_n\theta^n.$$

For example, if $n = 4$, the zero locus of the principal A-determinant (where the ML degree drops from four) is cut out by the discriminant of a degree four polynomial. The variety

where the ML degree drops at least two consists of two irreducible components. The first one corresponds to polynomials of the form $f(\theta) = k(\theta + a)^2(\theta + b)^2$. The variety where the ML degree is one is a determinantal variety, defined by 2-minors of the matrix

$$\begin{pmatrix} c_1 & 4c_2 & 3c_3 & 8c_4 \\ 2c_0 & 3c_1 & c_2 & c_3 \end{pmatrix}.$$

Example 4.3. In the case that $k = 2$, the toric variety is known as a *Hirzebruch surface*, and we can draw the corresponding polytope (Figure 2). We denote this surface by $\mathcal{H}_{n_1,n_2}$.

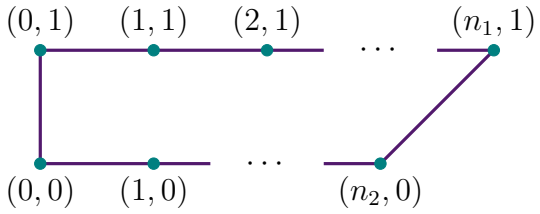

FIGURE 2. Polytope in $\mathbb{R}^2$ corresponding to a Hirzebruch surface.

For $c = (1, 1, \ldots, 1, 1)$, the ML degree of $\mathcal{H}_{n_1,n_2}$ is $n_1 + n_2 - \gcd(n_1 + 1, n_2 + 1) + 1$. This is because $g_1 = 1 + \theta + \cdots + \theta^{n_1}$ and $g_2 = 1 + \theta + \cdots + \theta^{n_2}$. The roots of g_1 and g_2 are respectively the $(n_1 + 1)$th and $(n_2 + 1)$th roots of unity not equal to 1. It follows that there are $\gcd(n_1 + 1, n_2 + 1) - 1$ repeated roots of $g_1 g_2$.

Example 4.4. If we choose the coefficients c_{ij} to be the binomial coefficients $\binom{n_i}{j}$, then $g_i = (1 + \theta_{d-1})^{n_i}$ for all i. For this choice of coefficients, the ML degree is 1 by Theorem 4.1, and the likelihood equations (2) have one solution given by

$$\theta_i = \frac{-u_{(d-1)+}u_{i+}}{u_{++}^2}(1 + \theta_{d-1})^{n_{d-1}-n_i} \text{ for } i = 1, 2, \ldots, d-2,$$

where $u_{++} = \sum_i \sum_j u_{ij}$, and θ_{d-1} equals the unique solution to

$$(Au)_{d-1}(1 + \theta_{d-1}) = \theta_{d-1}(u_{(d-1)+} + \sum_{i=1}^{d-2} n_i u_{i+}).$$

5. Veronese Embeddings

In this section we study Veronese and Veronese-type varieties.

Definition 5.1. Consider the $(d-1) \times n$ integer matrix A with columns that are the non-negative integer vectors whose coordinate sum $\leq k$. The projective toric variety defined by A is the *Veronese* $\mathrm{Ver}(d-1, k)$ for $d, k \geq 1$.

It can be seen that $\deg(\mathrm{Ver}(d-1, k)) = k^{d-1}$. A conjecture presented in [Vuc16] is that under the standard embedding, that is, with scaling vector $c = (1, 1, \ldots, 1, 1)$, the ML degree of any Veronese variety equals its degree. We now show this is true when $k \leq d - 1$.

Proposition 5.2. *Consider* $\mathrm{Ver}(d-1, k)$ *for* $k \leq d-1$ *which is embedded using the map* (1) *with scaling given by* $c = (1, 1, \ldots, 1, 1)$. *Then* $c \notin \Sigma_A$ *and* $\mathrm{mldeg}(\mathrm{Ver}(d-1, k)) = k^{d-1}$.

Proof. First note that, by induction, it is sufficient to show that $c \notin \Delta_A$. To prove this we need to show that the polynomial

$$f = \sum \theta^{a_i} = 1 + \sum_{i=1}^{k} h_i$$

is nonsingular, where h_i denotes the ith complete homogeneous symmetric polynomials in $\theta_1, \ldots, \theta_{d-1}$. The complete homogeneous symmetric polynomial $h_1, \ldots, h_{d-1}$ form an algebra basis for Λ_{d-1}, the ring of symmetric polynomials in variables $\theta_1, \ldots, \theta_{d-1}$, see [Mac98, (2.8)]. The regular embedding φ of Λ_{d-1} into the polynomial ring $\mathbb{C}[\theta_1, \ldots, \theta_{d-1}]$ (specified by the definition of h_i in terms of θ_j's) induces an embedding φ^* of $\mathrm{Spec}(\Lambda_{d-1})$ into $\mathbb{A}^{d-1} = \mathrm{Spec}(\mathbb{C}[\theta_1, \ldots, \theta_{d-1}])$. Consider the hypersurface $V = V(1 + h_1 + h_2 + \cdots + h_k) \subset \mathrm{Spec}(\Lambda_{d-1})$. For $k \leq d-1$ we have that $1 + h_1 + h_2 + \cdots + h_k$ is a linear form in the indeterminates $h_1, \ldots, h_{d-1}$. Hence V is smooth in $\mathrm{Spec}(\Lambda_{d-1})$. Since the induced map φ^* is a regular embedding, $\varphi^*(V)$ is smooth in $\mathbb{A}^{d-1}$, and by the definition of h_i and the map φ we have that $\varphi(1 + \sum_{i=1}^{k} h_i) = f$.Thus f defines a smooth hypersurface in $\mathbb{A}^{d-1}$, and $c = (1, 1, \ldots, 1, 1) \notin \Delta_A$. The conclusion follows by Theorem 3.1. □

Remark 5.3 (Hypersimplex). Let $k \leq d-1$. Consider the $(d-1) \times n$ integer matrix A where $n = \binom{d}{k}$ and where the columns of A are the vectors in $\{0,1\}^{d-1}$ that have precisely k or $k-1$ entries equal to 1. The $(d-1)$-dimensional polytope $Q = \mathrm{conv}(A)$ is called the *hypersimplex*. The projective toric variety V associated to A represents generic torus orbits on the Grassmannian of k-dimensional linear subspaces in $\mathbb{C}^d$. It is shown in the discussion preceding [HS17, Proposition 4.7] that $c = (1, 1, \ldots, 1, 1) \notin \Sigma_A$ for this matrix A. We remark that this result can be seen from the same proof given in Proposition 5.2.

Let e_{k-1} and e_k be elementary symmetric polynomials in $\theta_1, \ldots, \theta_{d-1}$ of degree $k-1$ and k, respectively. Similar to above, we need to only show that the polynomial

$$f = e_{k-1} + e_k$$

is nonsingular. This proceeds in a manner identical to the proof of Proposition 5.2 since the elementary symmetric polynomials also form an algebra basis for Λ_{d-1}, and hence are algebraically independent (see [Mac98, (2.4)]). Thus, for the projective toric variety V associated to the hypersimplex, $\mathrm{mldeg}(V)$ is equal to the normalized volume of the hypersimplex Q, which is the Eulerian number $A(d-1, k-1)$.

Returning to $\mathrm{Ver}(d-1, k)$, we note that in the case $k = 2$, the polynomial f reduces to a quadratic form, and a direct proof that $c = (1, 1, \ldots, 1, 1) \notin \Sigma_A$ can be given by looking at the principal minors of the corresponding symmetric matrix

$$C = \begin{pmatrix} 2c_{00} & c_{01} & \cdots & c_{0(d-1)} \\ c_{01} & 2c_{11} & \cdots & c_{1(d-1)} \\ \vdots & \vdots & \ddots & \vdots \\ c_{0(d-1)} & c_{1(d-1)} & \cdots & 2c_{(d-1)(d-1)} \end{pmatrix}, \tag{8}$$

where $f = (1, \theta_1, \ldots, \theta_{d-1}) C (1, \theta_1, \ldots, \theta_{d-1})^T$.

Proposition 5.4. *Consider* $\mathrm{Ver}(d-1, 2)$. *Then* $c = (1, 1, \ldots, 1, 1) \notin \Sigma_A$. *In particular, we have that* $\mathrm{mldeg}(\mathrm{Ver}(d-1, 2)) = \deg(\mathrm{Ver}(d-1, 2)) = 2^{d-1}$.

Proof. Let $Q = \text{conv}(A)$. From (5) we have that

$$\Sigma_A = \bigcup_{\Gamma \text{ face of } Q} \nabla_{A\cap\Gamma}.$$

The A-discriminant of the toric variety corresponding to each face will be given by a principal minor (of size corresponding to the dimension of the face) of the symmetric matrix C in (8). Any principal minor of size $r \times r$ will have the same form as C. Hence, to show that $c = (1, 1, \ldots, 1, 1) \notin \Sigma_A$ we need only verify that the $r \times r$ matrix

$$\tilde{C} = C(1, 1, \ldots, 1, 1) = \begin{pmatrix} 2 & 1 & \cdots & 1 \\ 1 & 2 & \cdots & \vdots \\ \vdots & 1 & \ddots & 1 \\ 1 & 1 & \cdots & 2 \end{pmatrix}$$

has a nonzero determinant. A calculation shows that $\det(\tilde{C}) = r + 1 \neq 0$ for all $r \geq 2$. Thus, by Theorem 3.1, it follows that $\text{mldeg}(\text{Ver}(d-1, 2)) = \deg(\text{Ver}(d-1, 2)) = 2^{d-1}$. □

Now we give a sufficient condition for the opposite effect to happen: instead of having the ML degree equal to the degree, the ML degree equals to one. The polynomial f has terms of degree $\leq k$ in $\theta_1, \ldots, \theta_{d-1}$. It uniquely corresponds to a symmetric tensor C. We say C has tensor rank one if the polynomial f is a power of a linear form.

Theorem 5.5. *Let C be the symmetric tensor corresponding to f where*

$$f = \sum_{\deg(\theta^a)\leq k} c_a \theta^a.$$

If the tensor rank of C is equal to one then $\text{mldeg}(\text{Ver}(d-1, k)^c) = 1$.

Proof. In this case we have that $f = L^k$ where $L = b_0 + b_1\theta_1 + \cdots + b_{d-1}\theta_{d-1}$. Then in the likelihood equations (3) we have that $L \neq 0$ and

$$\begin{array}{rcl} (Au)_1 L^k & = & (u_+ k b_1)\theta_1 L^{k-1} \\ (Au)_2 L^k & = & (u_+ k b_2)\theta_2 L^{k-1} \\ & \vdots & \\ (Au)_{d-1} L^k & = & (u_+ k b_{d-1})\theta_{d-1} L^{k-1}. \end{array} \tag{9}$$

So the factor L^{k-1} cancels out and we obtain a linear system in the θ_i's. □

Example 5.6. Let $d = 3$ and $k = 2$, that is, we consider the second Veronese embedding of $\mathbb{P}^2$. In this case the matrix defining $\text{Ver}(2, 2)$ is $A = \begin{pmatrix} 0 & 1 & 2 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 2 & 1 \end{pmatrix}$. The polytope $Q = \text{conv}(A)$ is given in Figure 3 with the labeled lattice point a_i corresponding to the ith column of the matrix A.

By (5), the ideal of $\Sigma_A \subset (\mathbb{C}^*)^6$ is defined by the polynomial E_A in (10),

$$\begin{aligned} E_A &= \Delta_A \cdot \Delta_{[a_1\ a_4\ a_5]} \cdot \Delta_{[a_3\ a_5\ a_6]} \cdot \Delta_{[a_1\ a_2\ a_3]} \\ &= \det \begin{pmatrix} 2c_{00} & c_{10} & c_{01} \\ c_{10} & 2c_{20} & c_{11} \\ c_{01} & c_{11} & 2c_{02} \end{pmatrix} \det \begin{pmatrix} 2c_{00} & c_{10} \\ c_{10} & 2c_{20} \end{pmatrix} \det \begin{pmatrix} 2c_{20} & c_{11} \\ c_{11} & 2c_{02} \end{pmatrix} \det \begin{pmatrix} 2c_{00} & c_{01} \\ c_{01} & 2c_{02} \end{pmatrix}. \end{aligned} \tag{10}$$

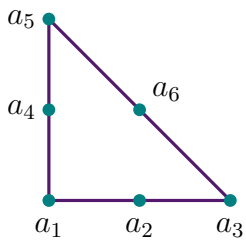

FIGURE 3. The polytope $Q = \text{conv}(A)$ of $\text{Ver}(2,2)$.

For generic values of c_{ij} in C as in equation (8), including $c_{ij} = 1$, we have that $\text{mldeg}(V^c) = 4 = 2^2$. On the other hand, if we take a scaling c so that C is the matrix with all entries equal to 2, the ML degree drops to $\text{mldeg}(V^c) = 1$. See Table 2 for other combinations. Note that when the rank of C is 1, the ML degree is 1, illustrating Theorem 5.5. We also see that the converse does not hold: the ML degree can be 1 even with the rank of C equal to 3.

C	Δ_A	$\Delta_{[a_1\ a_4\ a_5]}$	$\Delta_{[a_3\ a_5\ a_6]}$	$\Delta_{[a_1\ a_2\ a_2]}$	mldeg
$\begin{bmatrix} 2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{bmatrix}$	$\neq 0$	$\neq 0$	$\neq 0$	$\neq 0$	**4**
$\begin{bmatrix} 2 & 2 & 1 \\ 2 & 2 & 3 \\ 1 & 3 & 2 \end{bmatrix}$	0	0	$\neq 0$	$\neq 0$	**3**
$\begin{bmatrix} 2 & 2 & 1 \\ 2 & 2 & 2 \\ 1 & 2 & 2 \end{bmatrix}$	0	0	0	$\neq 0$	**2**
$\begin{bmatrix} -2 & 2 & 2 \\ 2 & -2 & 2 \\ 2 & 2 & -2 \end{bmatrix}$	$\neq 0$	0	0	0	**1**
$\begin{bmatrix} 17 & 22 & 27 \\ 22 & 29 & 36 \\ 27 & 36 & 45 \end{bmatrix}$	0	$\neq 0$	$\neq 0$	$\neq 0$	**3**
$\begin{bmatrix} 2 & 3 & 3 \\ 3 & 5 & 5 \\ 3 & 5 & 5 \end{bmatrix}$	0	$\neq 0$	0	$\neq 0$	**2**
$\begin{bmatrix} 2 & 2 & 2 \\ 2 & 2 & 2 \\ 2 & 2 & 2 \end{bmatrix}$	0	0	0	0	**1**

TABLE 2. The ML degree of $(\text{Ver}(2,2)^c)$ for different scalings c_{ij} in the matrix C.

6. Segre Embeddings

In this section, we study the maximum likelihood degree of the scaled Segre embedding V^c. We give a sufficient condition for V^c to have ML degree one.

For $c \in (\mathbb{C}^*)^{mn}$, let ψ^c be the map defined by

$$\psi^c\left(s, \theta_1^{(1)}, \ldots, \theta_m^{(1)}, \theta_1^{(2)}, \ldots, \theta_n^{(2)}\right) = \left(c_{11} s\theta_1^{(1)}\theta_1^{(2)}, \ldots, c_{ij} s\theta_i^{(1)}\theta_j^{(2)}, \ldots, c_{mn} s\theta_m^{(1)}\theta_n^{(2)}\right) \tag{11}$$

Then V^c is a scaled Segre embedding of $\mathbb{P}^{m-1} \times \mathbb{P}^{n-1}$. In terms of the coordinates p_{ij} with $1 \le i \le m$ and $1 \le j \le n$ on $\mathbb{P}^{mn-1}$, the defining ideal I^c of V^c is given by the 2-minors of the $m \times n$ matrix with ijth entry $\frac{p_{ij}}{c_{ij}}$.

From this it follows that the likelihood equations (4) are

$$\frac{u_{i+}}{u_{++}} = \frac{p_{i+}}{p_{++}} \text{ for } 1 \le i \le m, \quad \frac{u_{+j}}{u_{++}} = \frac{p_{+j}}{p_{++}} \text{ for } 1 \le j \le n, \quad \frac{p_{ij}p_{kl}}{c_{ij}c_{kl}} = \frac{p_{il}p_{kj}}{c_{il}c_{kj}}.$$

Here $u_{i+} = \sum_{j=1}^n u_{ij}$, $u_{+j} = \sum_{i=1}^m u_{ij}$, and $u_{++} = \sum_{i=1}^m \sum_{j=1}^n u_{ij}$. We define p_{i+}, p_{+j}, and p_{++} analogously. [Example 1.12]PS05 and [DSS09, Example 2.1.2]).

For all $c \in (\mathbb{C}^*)^{mn}$, we have $\deg(V^c) = \binom{m+n-2}{m-1}$.

If $c = (1, 1, \ldots, 1, 1)$, then the ML degree of V^c is well known to be one (see [PS05, Example 1.12] and [DSS09, Example 2.1.2]). Also the principal A-determinant for Segre embedding is easy to describe [GKZ94, Chapter 9, Section 1]:

$$E_A = \prod_{i=1}^{\min(m,n)} \prod_{\substack{1 \le a_1 < \cdots < a_i \le m \\ 1 \le b_1 < \cdots < b_i \le n}} \det[a_1, \ldots, a_i; b_1, \ldots, b_i] \tag{12}$$

where $[a_1, \ldots, a_i; b_1, \ldots, b_i]$ is the submatrix of c indexed by the rows $1 \le a_1 < \cdots < a_i \le m$ and the columns $1 \le b_1 < \cdots < b_i \le n$.

Example 6.1. Let V^c be the scaled Segre variety with scaling

$$c = \begin{pmatrix} 1 & c_{12} & c_{13} \\ c_{21} & c_{22} & 1 \\ 1 & 1 & 1 \end{pmatrix} \in M_{3,3}(\mathbb{C}^*).$$

Here, $m = n = 3$, and the matrix A is given by

$$A = \begin{pmatrix} 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 \end{pmatrix}.$$

The degree of V^c is $\binom{3+3-2}{3-1} = 6$. The assignments

$$(c_{12}, c_{13}, c_{21}, c_{22}) \in \{(1,1,1,1),\ (2,1,1,1),\ (2,3,1,1),\ (2,3,1,2),\ (2,3,2,1),\ (2,3,2,3)\}$$

produce, respectively, the ML degrees $\mathrm{mldeg}(V^c) = 1, 2, 3, 4, 5, 6$. This shows V^c can have any ML degree i, for $1 \le i \le \deg(V^c)$.

Conjecture 6.2. *For each $i \in \{1, \ldots, \deg(V)\}$, there exists $c \in (\mathbb{C}^*)^{mn}$ such that the scaled Segre variety V^c has ML degree i.*

Example 6.3. The scaled toric variety V^c with the configuration

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \end{pmatrix}$$

is isomorphic to the scaled Segre embedding of $\mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1 \times \mathbb{P}^1$ into $\mathbb{P}^{15}$. The degree of V^c in this instance is 24. For each of the possible ML degrees $i = 1, 2, \ldots, 24$, we have found a scaling vector $c \in (\mathbb{C}^*)^{16}$ for which $\mathrm{mldeg}(V^c) = i$. We see from (12) that for a point to lay in the vanishing of the principal A-determinant we must have that certain minors of the matrix (c_{ij}) vanish; this observation informed our choice of scaling vector c.

We conclude this section with a result analogous to Theorem 5.5.

Proposition 6.4. *If* $\mathrm{rank}(c) = 1$, *then* $\mathrm{mldeg}(V^c) = 1$.

Proof. If $\mathrm{rank}(c) = 1$ then all the 2-minors of c vanish. This implies that $c_{ij}c_{kl} = c_{il}c_{kj}$ for all $1 \le i, k \le m$ and $1 \le j, l \le n$. Then our likelihood equations reduce to

$$\frac{u_{i+}}{u_{++}} = \frac{p_{i+}}{1} \text{ for } 1 \le i \le m \quad \frac{u_{+j}}{u_{++}} = \frac{p_{+j}}{1} \text{ for } 1 \le j \le n, \quad p_{ij}p_{kl} = p_{il}p_{kj},$$

which are the likelihood equations for the unscaled Segre embedding of $\mathbb{P}^{m-1} \times \mathbb{P}^{n-1}$. This is known to have ML degree equal to one. □

7. Hierarchical loglinear and graphical models

A large class of discrete exponential models that are used in statistical practice are hierarchical loglinear models [BFH07] and undirected graphical models [Lau96]. In both cases, the toric variety is constructed via a parametrization based on a simplicial complex where the vertices of the simplicial complex correspond to discrete random variables. In the second case, the simplicial complex is the clique complex of a given graph. These toric varieties, their defining ideals, and the polyhedral geometry of $Q = \mathrm{conv}(A)$ of the underlying matrix A have been intensely studied starting with [HS02] for the hierarchical loglinear models and [GMS06] for graphical models.

Decomposable models form a subclass of these models, and they present attractive factorization properties. For instance, a graphical model is decomposable if and only if the underlying graph is a chordal graph [Lau96]. In particular, we have the following.

Theorem 7.1. [GMS06, Theorem 4.4] *A hierarchical loglinear model, and hence an undirected graphical model, is decomposable if and only if the ML degree of the corresponding toric variety is one.*

For example, the (unscaled) Segre embedding of $\mathbb{P}^{m_1} \times \cdots \times \mathbb{P}^{m_k}$ is a decomposable graphical model, and hence has ML degree one. In this section we consider some examples of nondecomposable models. First we study a family of examples known as no-three-way interaction models.

7.1. No-three-way interaction models.

Example 7.2 (Binary 3-cycle)**.** Consider the hierarchical loglinear model arising from the 3-cycle depicted in Figure 4. This is also known as the no-three-way interaction model. We

let the random variables X, Y, Z corresponding to the vertices to be binary variables. One parametrization $\psi : (\mathbb{C}^*)^6 \to (\mathbb{C}^*)^8$ is given by the matrix

$$A = \begin{pmatrix} 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 \end{pmatrix}.$$

We see that the 6×8 matrix A is in general position as in Section 3.1. Hence, the toric variety V defined by A is a hypersurface with a generator of full support. For $i, j, k \in \{0, 1\}$, let $p_{ijk} = \text{prob}(X = i, Y = j, Z = k)$. Labeling the columns of the matrix with the p_{ijk} in increasing lexicographic order, we find that the unique toric generator is

$$g = p_{000}p_{011}p_{101}p_{110} - p_{001}p_{010}p_{100}p_{111}.$$

By Corollary 3.4, the principal A-determinant is equal to the A-discriminant, and it is

$$\Delta_A = c_{000}c_{011}c_{101}c_{110} - c_{001}c_{010}c_{100}c_{111}.$$

The toric variety itself has degree 4 since it has a single degree 4 generator. However, if we choose the scaling coefficients $c_{ijk} = 1$ for all $i, j, k \in \{0, 1\}$, we see that this scaling vector lies in Σ_A, and thus $\text{mldeg}(V^c)$ drops, in this case, to 3. In fact, for this example, the only possibilities for $\text{mldeg}(V^c)$ are 3 when the scaling vector lies in Σ_A and 4 when the scaling vector does not.

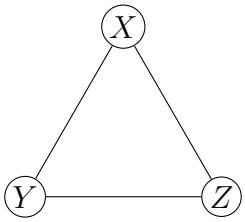

FIGURE 4. The 3-cycle.

Proposition 7.3. *The ML degree of the binary* 3*-cycle is* 4 *unless* $c \in (\mathbb{C}^*)^{d+1}$ *is in* Σ_A*. If* $c \in \Sigma_A$*, then* $\text{mldeg}(V^c) = 3$.

Proof. We fix the lexicographic monomial ordering

$$p_{000} > p_{001} > p_{010} > \cdots > p_{111} > c_{000} > c_{001} > \cdots > c_{111}$$

over $\mathbb{Q}[p_{000}, p_{001}, p_{010}, p_{011}, p_{100}, p_{101}, p_{110}, p_{111}, c_{000}, c_{001}, c_{010}, c_{011}, c_{100}, c_{101}, c_{110}, c_{111}]$.

Let $u = (u_{000}, \ldots, u_{111})$ be a data vector in $\mathbb{Z}^8$ and let $u_+ = \sum u_{ijk}$. Consider

$$I = \langle\, g\,,\, \Delta_A\,,\, L_f(p) - 1\,,\, u_+L_1(p) - (Au)_1\,,\, \ldots\,,\, u_+L_6(p) - (Au)_6\,\rangle$$

where $f = \sum c_{ijk} p_{ijk}$. Using [GS] and the random data vector $u = (2, 3, 5, 7, 11, 13, 17, 19)$, we calculate a Gröbner basis for I:

$$
\begin{aligned}
g_1 &= c_{000}c_{011}c_{101}c_{110} - c_{001}c_{010}c_{100}c_{111} \\
g_2 &= 5021863 p_{111}^3 c_{111}^3 - 3752210 p_{111}^2 c_{111}^2 + 984280 p_{111} c_{111} - 89856 \\
g_3 &= 77 p_{110} c_{110} + 77 p_{111} c_{111} - 36 \\
g_4 &= 17472 p_{110} c_{001} c_{010} c_{100} - 456533 p_{111}^3 c_{000} c_{011} c_{101} c_{111}^2 + 341110 p_{111}^2 c_{000} c_{011} c_{101} c_{111} + \cdots \\
g_5 &= 77 p_{101} c_{101} + 77 p_{111} c_{111} - 32 \\
g_6 &= 19656 p_{101} c_{001} c_{010} c_{100} - 456533 p_{111}^3 c_{000} c_{011} c_{110} c_{111}^2 + 341110 p_{111}^2 c_{000} c_{011} c_{110} c_{111} + \cdots \\
g_7 &= 77 p_{100} c_{100} - 77 p_{111} c_{111} + 8 \\
g_8 &= 184320 p_{100} c_{000} c_{011} - 73501813 p_{101} p_{110} p_{111}^2 c_{001} c_{010} c_{111}^3 + \cdots \\
g_9 &= 77 p_{011} c_{011} + 77 p_{111} c_{111} - 26 \\
g_{10} &= 359424 p_{011} c_{001} c_{010} - 115502849 p_{100} p_{111}^3 c_{000} c_{101} c_{110} c_{111}^2 + \cdots \\
g_{11} &= 11 p_{010} c_{010} - 11 p_{111} c_{111} + 2 \\
g_{12} &= 287539200 p_{010} c_{000} - 1097794317389 p_{011} p_{101} p_{110} p_{111}^2 c_{001} c_{100} c_{111}^3 + \cdots \\
g_{13} &= 77 p_{001} c_{001} - 77 p_{111} c_{111} + 16 \\
g_{14} &= 4313088 p_{001} c_{000} - 12619941719 p_{011} p_{101} p_{110} p_{111}^2 c_{010} c_{100} c_{111}^3 + \cdots \\
g_{15} &= 1725235200 p_{000} - 59243384844259 p_{001} p_{010} p_{100} p_{111}^3 c_{011} c_{101} c_{110} c_{111}^2 + \cdots
\end{aligned}
$$

As long as each $c_{ijk} \in \mathbb{C}^*$ and satisfies the equation $g_1 = \Delta_A = 0$, we see that g_2 is a univariate polynomial in p_{111} of degree 3, and the initial terms of g_3 through g_{15} have degree 1 in p_{ijk}. Therefore, when the coefficient vector lies in Σ_A, the ML degree is always 3. □

More generally, we consider the no-three-way interaction model C_m based on the 3-cycle with one m-ary variable and two binary variables. That is, for $i \in \{0, \ldots, m-1\}$ and $j, k \in \{0, 1\}$, we set

$$p_{ijk} = \text{prob}(X = i, Y = j, Z = k) = a_{ij} b_{ik} c_{jk}.$$

One can compute the normalized volume of $Q = \text{conv}(A_m)$ for the matrix A_m associated C_m. This gives a formula for the degree of C_m.

Proposition 7.4. *The degree of C_m is $m2^{m-1}$.*

When $m = 3$, $\deg(C_3) = 12$ and $\text{mldeg}(C_3) = 7$. The toric ideal is

$$\langle\, p_{100}p_{111}p_{201}p_{210} - p_{101}p_{110}p_{200}p_{211},\ p_{000}p_{011}p_{101}p_{110} - p_{001}p_{010}p_{100}p_{111},\ p_{000}p_{011}p_{201}p_{210} - p_{001}p_{010}p_{200}p_{211} \,\rangle.$$

Observe that the toric variety is not a hypersurface, and thus calculating the principal A-determinant requires more work than the previous example. When we compute Δ_{A_3}, the A-discriminant for the whole polytope $Q = \text{conv}(A_3)$, we do *not* get a hypersurface:

$$
\begin{aligned}
\Delta_A = \langle & c_{100}c_{111}c_{201}c_{210} - c_{101}c_{110}c_{200}c_{211}, \\
& c_{000}c_{011}c_{101}c_{110} - c_{001}c_{010}c_{100}c_{111}, \\
& c_{000}c_{011}c_{201}c_{210} - c_{001}c_{010}c_{200}c_{211} \rangle.
\end{aligned}
$$

We already see that $c = (1, 1, \ldots, 1, 1)$ is in Σ_A and hence we know $\text{mldeg}(C_3^c) < \deg(C_3) = 12$. Also comparing the toric variety and the discriminant locus ∇_{A_3} we see that they are identical. Such a toric variety is known as self-dual. Based on our computations for larger m we state the following conjecture.

Conjecture 7.5. *The no-three-way interaction model with one m-ary variable and two binary variables is self-dual.*

To compute the entire principal A-determinant, we need to consider all faces of Q. If a face Γ has dimension e and $|\Gamma \cap A_3| = e + 1$, then $\Delta_{\Gamma \cap A_3}$ is the unit ideal. Moreover, for all codimension one and codimension two faces Γ which are not simplices, we observe that $\nabla_{\Gamma \cap A_3}$ are not hypersurfaces. They also lie in a coordinate hyperplane. For instance, for the facet Γ where the elements of $\Gamma \cap A_3$ correspond to

$$\{p_{001}, p_{011}, p_{100}, p_{101}, p_{110}, p_{111}, p_{200}, p_{201}, p_{210}, p_{211}\}$$

the discriminant is defined by

$$\langle\, c_{001},\ c_{011},\ c_{100}c_{111}c_{201}c_{210} - c_{101}c_{110}c_{200}c_{211} \,\rangle.$$

Therefore, for $c = (1, 1, \ldots, 1, 1)$, these discriminants do not contribute to a drop in the ML degree. There is a total of three codimension 3 faces that are not simplices. In each case $\Gamma \cap A_3$ is in general position and their discriminants are given by Corollary 3.4. We list the indeterminates which correspond to the elements of $\Gamma \cap A_3$ in each of these three faces:

$$\{p_{000}, p_{001}, p_{010}, p_{011}, p_{100}, p_{101}, p_{110}, p_{111}\},$$
$$\{p_{000}, p_{001}, p_{010}, p_{011}, p_{200}, p_{201}, p_{210}, p_{211}\},$$
$$\{p_{100}, p_{101}, p_{110}, p_{111}, p_{200}, p_{201}, p_{210}, p_{211}\}.$$

It turns out that each face of codimension > 3 is a simplex and there are no more discriminants contributing to the principal A-determinant. Based on our computations for larger m, we state the following conjecture.

Conjecture 7.6. *The ML degree of C_m is $2^m - 1$.*

7.2. **The binary 4-cycle.** The binary 4-cycle is the model we have used in Example 1.2. Let S, B, H, L be four binary random variables and let $X = (S, B, H, L)$ be the joint random variable where we set

$$p_{ijk\ell} \ = \ \text{prob}(S = i, B = j, H = k, L = \ell).$$

The family of probability distributions for X which factor according to the graphical model depicted in Figure 1 can be described by the following monomial parametrization: let $a_{ij}, b_{jk}, c_{k\ell}, d_{i\ell}$ be parameters for $i, j, k, \ell \in \{0, 1\}$ and let $p_{ijk\ell} = a_{ij}b_{jk}c_{k\ell}d_{i\ell}$. Below is the matrix A arising from a different parametrization that gives a full rank matrix. The columns are labeled by $p_{0000}, p_{0001}, \ldots, p_{1111}$ in increasing lexicographic ordering.

$$A = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \end{pmatrix}$$

The Zariski closure V of the image of this parametrization is a toric variety that is defined by the following prime ideal:

$$\begin{aligned} I = \langle\, & p_{1011}p_{1110} - p_{1010}p_{1111},\ p_{0111}p_{1101} - p_{0101}p_{1111},\ p_{1001}p_{1100} - p_{1000}p_{1101},\ p_{0110}p_{1100} - p_{0100}p_{1110}, \\ & p_{0011}p_{1001} - p_{0001}p_{1011},\ p_{0011}p_{0110} - p_{0010}p_{0111},\ p_{0001}p_{0100} - p_{0000}p_{0101},\ p_{0010}p_{1000} - p_{0000}p_{1010}, \\ & p_{0100}p_{0111}p_{1001}p_{1010} - p_{0101}p_{0110}p_{1000}p_{1011},\ p_{0010}p_{0101}p_{1011}p_{1100} - p_{0011}p_{0100}p_{1010}p_{1101}, \\ & p_{0001}p_{0110}p_{1010}p_{1101} - p_{0010}p_{0101}p_{1001}p_{1110},\ p_{0001}p_{0111}p_{1010}p_{1100} - p_{0011}p_{0101}p_{1000}p_{1110}, \\ & p_{0000}p_{0011}p_{1101}p_{1110} - p_{0001}p_{0010}p_{1100}p_{1111},\ p_{0000}p_{0111}p_{1001}p_{1110} - p_{0001}p_{0110}p_{1000}p_{1111}, \\ & p_{0000}p_{0111}p_{1011}p_{1100} - p_{0011}p_{0100}p_{1000}p_{1111},\ p_{0000}p_{0110}p_{1011}p_{1101} - p_{0010}p_{0100}p_{1001}p_{1111}\,\rangle. \end{aligned}$$

The degree of this toric variety is 64 and $\mathrm{mldeg}(V) = 13$. This was computed in [GMS06, p. 1484] where the question of explaining the fact $\mathrm{mldeg}(V) = 13$ was first raised. One road to an explanation is to compute the A-discriminants of all the faces of $Q = \mathrm{conv}(A)$, and then determine the contribution of each for the drop in the ML degree. There are multiple faces that might contribute to such a drop. At this moment, we do not understand how different discriminants *interact*. Also, it is not possible to compute every discriminant. For instance, with standard elimination methods we were not able to compute Δ_A corresponding to the entire polytope. However, one can check that the polynomial $f = \sum_i^{16} \theta^{a_i}$ where $c = (1, 1, \ldots, 1, 1)$ has singularities in $(\mathbb{C}^*)^{d-1}$. Hence, we conclude that Δ_A does contribute to the drop from the generic ML degree. Moreover, the polytope Q has 24 facets, of which 16 have nontrivial discriminants. Again, we were not able to compute these discriminants.

There are 168 codimension two faces of Q, and only 104 have nontrivial discriminants. These we can classify with respect to $|\Gamma \cap A|$, the cardinality of elements of A on the face Γ. The possibilities for this cardinality are 8, 9, and 10:

a) There are 40 codimension two faces with 8 vertices that have nontrivial discriminants. Only 5 of them are hypersurfaces generated by polynomials such as

$$c_{0100}c_{1010}c_{0111}c_{1001} - c_{1000}c_{0101}c_{0110}c_{1011}.$$

The rest are not hypersurfaces, and they lie in coordinate subspaces. They are defined by ideals of the form

$$\langle\, c_{1001},\ c_{0111},\ c_{0110},\ c_{1000},\ c_{0100}c_{0001} - c_{0000}c_{1010}\,\rangle.$$

b) There are 32 codimension two faces with 9 vertices that have nontrivial discriminants. None of them are hypersurfaces, and they all lie in coordinate subspaces. They are all of the form

$$\langle\, c_{1001},\ c_{0110},\ c_{0100}c_{1010}c_{0001} + c_{1000}c_{0010}c_{0101} - c_{0000}c_{1010}c_{0001}\,\rangle.$$

c) There are 32 codimension two faces with 10 vertices that have nontrivial discriminants. They are all hypersurfaces and take the form

$$\langle\, c_{0100}c_{1010}c_{0011}c_{1001} - c_{0100}c_{1010}c_{0001}c_{1011} - c_{1000}c_{0100}c_{0101}c_{1011} + c_{0000}c_{1010}c_{0101}c_{1011}\,\rangle.$$

8. MLE Homotopies

In this section we use homotopy continuation to track between ML estimates of different scalings of a given statistical model. Moreover, the endpoints may correspond to different scalings of the model with different ML degrees.

Consider the case where we have found a particular scaling c_{easy} such that the ML degree drops to one, but we really wish to compute the ML estimate for the natural statistical model corresponding to the scaling c_{stat}. The strategy is to apply a parameter homotopy [SW05] between c_{easy} and c_{stat}, tracking the unique solution of the former to a solution of the latter. We argue that the endpoint is the unique ML estimate in Birch's Theorem 2.6.

Example 8.1. (Veronese) We illustrate the strategy with a Veronese model. Let

$$A = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 2 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 2 \end{pmatrix},$$

with $c_{stat} = (1,1,1,1,1,1)$ and $c_{easy} = (1,2,1,2,2,1)$. We can check (see Section 5) that the ML degree corresponding to c_{stat} is 4, while the ML degree for c_{easy} is 1.

Suppose we observe the data vector $u = (1,3,5,7,9,2)$. Computing the unique solution for c_{easy} we obtain the ML estimate $\hat{\theta}_{win} = (0.0493, 1.8333, 1.6667)$. We track this point with a parameter homotopy towards c_{stat}, and obtain the point $\hat{\theta}_{track} = (0.0863, 1.6326, 1.5150)$ with corresponding $\hat{p}_{stat} = (0.09, 0.14, 0.23, 0.13, 0.21, 0.20)$.

To verify this, we solve the critical equations for c_{stat} to obtain the four solutions

$$\begin{array}{cc} (0.2888, 1.4316, -1.8931), & (0.3039, -1.8847, 1.3470), \\ (0.8578, -0.7629, -0.7189), & (0.0863, 1.6326, 1.5150), \end{array}$$

where $\hat{\theta}_{stat}$ is the solution with positive coordinates. Observe that $\hat{\theta}_{track} = \hat{\theta}_{stat}$, as desired.

Let c_{easy} and c_{stat} be scalings with positive entries. Let $F_{easy}(u,\theta)$ be the difference of the left and right sides of the equations of Definition 2.3 with c taken to be c_{easy}, and similarly for $F_{stat}(u,\theta)$. We may now state the main result of this section.

Theorem 8.2. *Fix a generic data vector u with positive entries. Let $\hat{p}_{easy}$ and $\hat{p}_{stat}$ be the respective ML estimates for these scaling problems for the data u and consider the homotopy below:*

$$H(\theta, t) = t \cdot F_{easy}(u,\theta) + (1-t) \cdot F_{stat}(u,\theta).$$

Let γ denote the path of the homotopy whose start point (at $t = 1$) corresponds to $\hat{p}_{easy}$. Then the endpoint of γ (at $t = 0$) is $\hat{p}_{stat}$.

Proof. By Birch's theorem (Theorem 2.6), the likelihood equations are given by $A\hat{p} = \frac{1}{u_+} Au$ for a data vector u. Let $\hat{p}_{win}$ and $\hat{p}_{stat}$ be the two monomial vectors for c_{easy} and c_{stat} respectively. The homotopy above can be rewritten as $A \cdot (t\hat{p}_{c(t)} + \frac{1}{u_+} u)$ where $c(t) = tc_{stat} + (1-t)c_{easy}$. Since $c(t)$ is positive for any positive real c_{easy}, c_{stat} and $t \in [0,1]$, Birch's theorem states that there is exactly one positive real solution to this system at every point along the homotopy path. Thus, as long as no paths intersect, the unique solution for c_{easy} will track to the unique positive real solution for c_{stat}. But paths only intersect where the Jacobian of the system drops rank, which we now show cannot occur.

For a value of $\hat{p}$ at some point on the path, the Jacobian matrix can be written as the product

$$A \begin{pmatrix} \frac{\partial p_1}{\partial \theta_1} & \cdots & \frac{\partial p_1}{\partial \theta_d} \\ \vdots & \ddots & \vdots \\ \frac{\partial p_n}{\partial \theta_1} & \cdots & \frac{\partial p_n}{\partial \theta_d} \end{pmatrix}.$$

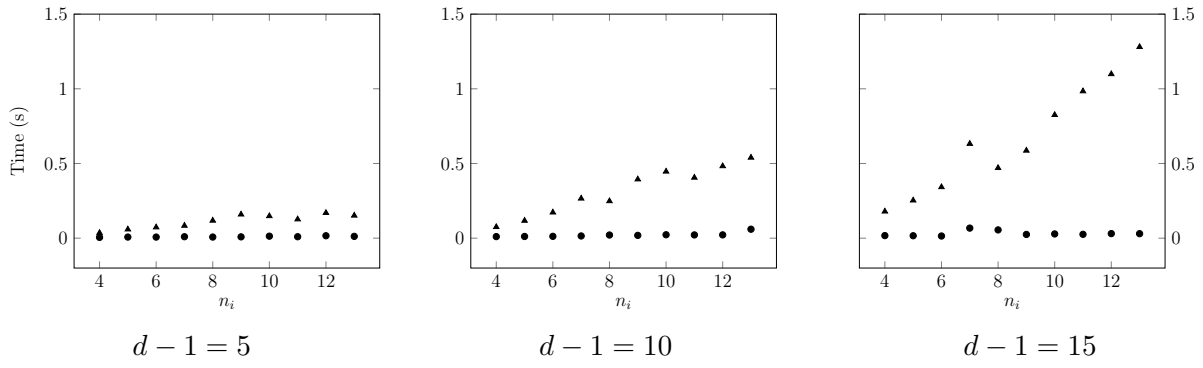

FIGURE 5. Running times of iterative proportional scaling (triangles) versus path tracking (circles) on the rational normal scroll (6) for three choices of $d-1$ with n_i's constant. Average of 7 trials.

Because we require the θ_i's to be nonzero, we can write this as

$$A \begin{pmatrix} \theta_1 \frac{\partial p_1}{\partial \theta_1} & \cdots & \theta_d \frac{\partial p_1}{\partial \theta_d} \\ \vdots & \ddots & \vdots \\ \theta_1 \frac{\partial p_n}{\partial \theta_1} & \cdots & \theta_d \frac{\partial p_n}{\partial \theta_d} \end{pmatrix} \operatorname{diag}\left(\frac{1}{\theta_1}, \ldots, \frac{1}{\theta_d}\right).$$

Note that if we denote $A = (a_{ij})$, then $\theta_i \frac{\partial p_j}{\partial \theta_i} = a_{ij} p_j$, so the middle matrix becomes $\operatorname{diag}(p_1, p_2, \ldots, p_n) A^T$ and the Jacobian thus factors as

$$A \operatorname{diag}(p_1, \ldots, p_n)\, A^T \operatorname{diag}\left(\frac{1}{\theta_1}, \ldots, \frac{1}{\theta_d}\right).$$

So the Jacobian is the product of these full rank matrices, hence has full rank. □

8.1. **Timings.** To test the viability of using homotopy methods, we ran timing comparisons with the iterative proportional scaling algorithm (IPS, Algorithm 2.1.9 in [DSS09]). We programmed IPS in `python` using `numpy` for fast linear algebra calculations and a target residual of $\epsilon = 10e^{-12}$. We choose the rational normal scroll (6) example from Section 4, setting $n_i = k$ for all i and varying $k \in \{4, \ldots, 13\}$ and $d - 1 \in \{5, 10, 15\}$. For the homotopy start system, we set $c_{ij} = \binom{n_i}{j}$ and tracked from the closed form solution given in Example 4.4. It is clear from Figure 5 that as the problem size grows, the homotopy method has a significant speed advantage over the iterative proportional scaling algorithm.[1]

REFERENCES

[BD76] Peter J. Bickel and Kjell A. Doksum, *Mathematical statistics*, Holden-Day, Inc., San Francisco, Calif.-Düsseldorf-Johannesburg, 1976, Basic ideas and selected topics, Holden-Day Series in Probability and Statistics.

[BFH07] Yvonne M. M. Bishop, Stephen E. Fienberg, and Paul W. Holland, *Discrete multivariate analysis: theory and practice*, Springer, New York, 2007, With the collaboration of Richard J. Light and Frederick Mosteller, Reprint of the 1975 original.

[CHKS06] Fabrizio Catanese, Serkan Hoşten, Amit Khetan, and Bernd Sturmfels, *The maximum likelihood degree*, Amer. J. Math. **128** (2006), no. 3, 671–697.

[CLS11] David A. Cox, John B. Little, and Henry K. Schenck, *Toric varieties*, Graduate Studies in Mathematics, vol. 124, American Mathematical Society, Providence, RI, 2011.

[1]Source code can be found at https://github.com/nbliss/Likelihood-MRC

[DR72] J. N. Darroch and D. Ratcliff, *Generalized iterative scaling for log-linear models*, Ann. Math. Statist. **43** (1972), 1470–1480.

[DSS09] Mathias Drton, Bernd Sturmfels, and Seth Sullivant, *Lectures on algebraic statistics*, Oberwolfach Seminars, vol. 39, Birkhäuser, Basel, Switzerland, 2009.

[EH85] David Edwards and Tomáš Havránek, *A fast procedure for model search in multidimensional contingency tables*, Biometrika **72** (1985), no. 2, 339–351. MR 801773

[Ful93] William Fulton, *Introduction to toric varieties*, Annals of Mathematics Studies, vol. 131, Princeton University Press, Princeton, NJ, 1993, The William H. Roever Lectures in Geometry.

[Ful98] ———, *Intersection theory*, 2nd ed., Springer, 1998.

[GKZ94] I. M. Gelfand, M. M. Kapranov, and A. V. Zelevinsky, *Discriminants, resultants, and multidimensional determinants*, Mathematics: Theory & Applications, Birkhäuser Boston, Inc., Boston, MA, 1994.

[GMS06] Dan Geiger, Christopher Meek, and Bernd Sturmfels, *On the toric algebra of graphical models*, Ann. Statist. **34** (2006), no. 3, 1463–1492.

[GS] Daniel R. Grayson and Michael E. Stillman, *Macaulay2, a software system for research in algebraic geometry*, Available at http://www.math.uiuc.edu/Macaulay2/.

[HKS05] Serkan Hoşten, Amit Khetan, and Bernd Sturmfels, *Solving the likelihood equations*, Found. Comput. Math. **5** (2005), no. 4, 389–407.

[HRS14] Jonathan Hauenstein, Jose Israel Rodriguez, and Bernd Sturmfels, *Maximum likelihood for matrices with rank constraints*, J. Algebr. Stat. **5** (2014), no. 1, 18–38.

[HS02] Serkan Hoşten and Seth Sullivant, *Gröbner bases and polyhedral geometry of reducible and cyclic models*, J. Combin. Theory Ser. A **100** (2002), no. 2, 277–301.

[HS14] June Huh and Bernd Sturmfels, *Likelihood geometry*, Combinatorial algebraic geometry, Lecture Notes in Math., vol. 2108, Springer, Cham, 2014, pp. 63–117.

[HS17] Martin Helmer and Bernd Sturmfels, *Nearest points on toric varieties*, Mathematica Scandinavica (2017), to appear.

[Huh13] June Huh, *The maximum likelihood degree of a very affine variety*, Compos. Math. **149** (2013), no. 8, 1245–1266.

[Huh14] June Huh, *Varieties with maximum likelihood degree one*, Journal of Algebraic Statistics **5** (2014), 1–17.

[Lau96] Steffen L. Lauritzen, *Graphical models*, Oxford Statistical Science Series, vol. 17, The Clarendon Press, Oxford University Press, New York, 1996, Oxford Science Publications.

[Mac98] Ian Grant Macdonald, *Symmetric functions and hall polynomials*, Oxford university press, 1998.

[Pet08] Sonja Petrović, *On the universal Gröbner bases of varieties of minimal degree*, Math. Res. Lett. **15** (2008), no. 6, 1211–1221.

[PS05] Lior Pachter and Bernd Sturmfels, *Algebraic statistics for computational biology*, Cambridge University Press, New York, NY, USA, 2005.

[R+81] Z. Reiniš et al., *Prognostic significance of the risk profile in the prevention of coronary heart disease*, Bratis. Lek. Listy **76** (1981), 137–150.

[SW05] Andrew Sommese and Charles Wampler, *The numerical solution of systems of polynomials arising in engineering and science*, World Scientific, Hackensack, NJ, USA, 2005.

[Vuc16] Radoslav Vuchkov, *The maximum likelihood degree of various toric varieties*, Master's thesis, San Francisco State University, 2016.